



xCoAx 2021 9th Conference on
Computation, Communication, Aesthetics & X

2021.xCoAx.org

Cyprus as AI Saw It: Digital Colonialism and AttnGAN Text to Image Synthesis

Keywords: GAN, Text to Image Synthesis, Cyprus, Digital Colonialism, Decolonisation

Marinos

Koutsomichalis

marinos@madlab.cool

Media Arts & Design Research
Lab, Cyprus University of
Technology, Limassol, Cyprus

Alexia Achilleos

alexia.a@madlab.cool

Media Arts & Design Research
Lab, Cyprus University of
Technology, Limassol, Cyprus

This paper discusses an ongoing project pivoting on an AttnGAN Text to Image synthesis as the means to critically comment on affairs of bias in historical and digital colonialism. It zooms in on Cyprus — a geographic region that is largely underrepresented in AI programming — as seen through the lenses of 19th century colonialism. A well-known AttnGAN pipeline is appropriated and trained over a DIY dataset to produce images that more genuinely reflect the (post-)colonial Cypriot reality — thus giving agency to a geographic region, the cultural and historical idiosyncrasies of which are generally not reflected in AI programming. Sample imagery is presented and the results of a qualitative survey are discussed in some detail. In this fashion, the authors intend to contribute empirical data to the ongoing discussion about digital colonialism, and to initiate a broader conceptualisation process about methodologies to decolonise current affairs in AI research. Technical, critical and ethical implications of the above are further discussed, with references to Digital Colonialism and AI-driven Text-to-Image synthesis.

1. Introduction and Motivation

This project draws upon a body of prior artistic experiments broadly taking a critical stance on the current state of affairs in AI-driven image synthesis, in an attempt to raise issues of historical and digital colonialism and zooming in on Cyprus — a geographic region that is grossly underrepresented in AI programming. To give a relevant example, consider *The New Riviera*,¹ wherein a pre-configured AttnGAN Text-to-Image (hence on, T2I) model is fed with real estate promotional text that targets foreign investors interested in the Cyprus citizenship investment program (via which one may easily acquire an EU passport). The resulting imagery, however, bears little, if any, resemblance to the Eastern Mediterranean landscape the input texts refer to; generated images are rather reminiscent of Western or North-European landscapes and architecture. Similarly, this endeavour zeroes in on Cyprus as seen through the eyes of a typical 19th century colonialist. The focus is on the *Cyprus, as I Saw it in 1879* book, wherein Sir Samuel Baker presents his impression on the British Empire's newly annexed colony (Baker 1879). Excerpts from the book are used to label training data, as well as the input to a T2I system that, as to be seen in a subsequent section, generates imagery that is expected to better resonate with Cypriot (post-)colonial state of affairs.

1. <https://alexiaachilleos.com/thenewriviera> March 2021

The overall motivation behind the project is to comment on digital colonialism related affairs in a hands-on critical design fashion. Accordingly, the authors intend to contribute empirical data to the broader discourse as well as to give agency to a geographic region, the cultural and historical idiosyncrasies of which are generally not reflected in present day AI research. In this fashion, the method discussed here aspires to become a paradigm of how questions of digital colonialism in AI-related research could be concretely raised in similar contexts. It should be underscored that the authors do not approach the local (post-)colonial reality from the outside — that is, they are not speaking on behalf of some oppressed other in some other part of the world. Being based in the East Mediterranean region they are themselves subjected to digital colonialism in its disparate manifestations.

Following this brief introductory section, the article proceeds with some background context on GAN-driven image synthesis and digital colonialism. The section after that is a technical overview of the T2I pipeline. Following, sample imagery is illustrated and the results of a qualitative survey concerning them are discussed. A Discussion section follows, wherein critical/ethical implications delineated. Concluding remarks and notes on future work sum up this paper.

2. Background

Following the emergence of AI systems that pivot on intelligent critics over the last few years, GAN-driven image synthesis has emerged as a prominent area of research with many relevant systems discussed in literature. Zooming on T2I pipelines, the first milestone has been achieved with Reed et al. (2016), wherein Natural Language Processing techniques are first used alongside the original GAN architecture introduced in Goodfellow et al. (2014). Improvements and variations of this approach are encountered in a number of other resources, such as in Zhang et al. (2017), discussing a ‘stacked’ T2I pipeline: the first stage GAN sketches the primitive shape and paint the colour of an object with respect to arbitrary text input, while the second stage employs a second GAN to rectify the latter and to amplify their resolution. The authors have further refined this system using multiple generators and multiple discriminators arranged in a tree-like structure to achieve better performance and to deliver even more persuasive photo-realistic imagery (Zhang et al., 2018). Zhu et al. also follow a multi-staged approach with a Dynamic Memory GAN model that allows the original text representation to change in between the different stages, in this way promising to deliver high-quality images even when the original ones are not very well generated (Zhu et al., 2019). Other models introduce Attention (AttnGAN), typically through a Deep Attentional Multimode Similarity Model (Show and Xu, 2015). In (Emir et al., 2019) we see an enhanced AttnGAN architecture featuring an attention module that performs Feature-wise Linear Modulation (FiLM) on text features. The authors argue that their model outperforms state-of-the-art methods in the generation of fashion-related imagery. With their MirrorGAN architecture, Qiao et al expand the AttnGAN model: here, images are generated from text with one GAN and then translated back to text via another GAN; cycle-consistency (a technique drawing on CycleGAN models) is then applied between the original and generated text descriptions (Qiao et al., 2019). A final example, of an altogether different approach, is (Dong et al., 2020), wherein the system is unsupervised and entirely relies on visual concepts that bridge two independent datasets comprising, respectively, image and text — no (labour-intensive to collect) pairwise image-text data are necessary here. The authors conclude that their model yields promising results that, occasionally, even outperform supervised pipelines.

Despite the abundance of working T2I pipelines, research of sorts primarily, if not exclusively, revolves around certain benchmarks that relate with the resolution of the resulting images, their fit for some particular task, or the overall performance of a model. That is to say that the subarea is largely governed by a

certain techno-solutionist/techno-fetishist paradigm. In reality, in all the above mentioned — as well as in the sheer majority of similar systems — the employed GANs are intentionally or unintentionally pre-configured to largely reflect a ‘Westerner’s gaze’ — one that implicitly favours the point of view of white, male, first-world, mid/high class, occidental perspective. This is, of course, well reflected in the training datasets that are readily available and most typically encountered in AI-driven image synthesis contexts. The field is currently almost exclusively governed by the logic of certain databases wherein non-western realities are scarcely — if at all — represented. Consider what is probably the three most important datasets at the moment: ‘The Oxford Flowers’ (Nilsback and Zisserman 2008), ‘The Caltech-UCSD Birds’ (hence on, CUB) (Wah et al., 2011), or the ‘Microsoft Common Objects in Context’ (hence on, COCO) (Lin et al., 2014) datasets. All of these former datasets have been created by major Western universities/organisations and, arguably, can be said to embody the present-day (digital) colonialist reality.

Accordingly, it can be taken for granted that the vast majority of existent T2I pipelines is not meant, and would in fact most often fail, to generate images accounting for non-western landscapes or for ones relating with other cultures even if fed with textual descriptions thereof. More than an assumption, this has been pragmatically confirmed to be the case in experimental settings such as, e.g., *The New Riviera* mentioned above, and is also demonstrated through example in Section 4 below. Western bias has been generally shown to surface AI-related systems of all disparate sorts and on many different occasions. A case worth referring to is that of the Beauty.AI AI-judged beauty contest that has been shown to significantly favour white individuals of Caucasian origin. Its creators have, indeed, acknowledged the inherent racism in their system and attribute it specifically to the fact that “*you might not have enough data, or the data might be biased*”.² That is, they point at the very real problem of most, if not all, existent datasets being colonised and Western-biased. In such a vein, this project is a very concrete attempt seeking to contribute non-biased DIY datasets that could be used to train AIs so that feminist and other intersectional ways of thinking can be better represented and made resonate across digital media (Sinders, 2019).

2. <https://www.theguardian.com/technology/2016/sep/08/artificial-intelligence-beauty-contest-doesnt-like-black-people>
March 2021

Beauty.AI is a rather indicative case of digital colonialism; that is, of practices that, deliberately or not, result in dominant cultures imposing their ideas upon others. Affairs of digital colonialism are actively accounted for and studied in a number of contexts; for instance, with respect to the impact of information and communication technologies on non-western knowledge systems (Young

2019), economic and political neo-imperialist practices (Kwet 2019), or third-world user-data exploitation (Coleman 2018). A critical perspective to digital colonialist practices suggests that there is a very real confrontation that oscillates between control and freedom and concerning individuals, populations, and entire geographic regions (Pinto 2019). This confrontation is simultaneously relevant at many different levels of cultural and political organisation and manifests by means of controlling digital assets, data, computational power, as well as by a variety of other means. Questions of sorts are very urging when it comes to AI-related research, of course; AI is already advertised as the flagship of a forthcoming new digital revolution and, as such, is expected to shape culture and governance in future societies. Even so, related technological infrastructure is being developed by a very small group of predominantly western countries, so that they end up ascribed with the symbolical, ethical and ideological values of the latter.

Benjamin (2019) Benjamin argues that algorithms do often accelerate hierarchies and do sustain or impose social divisions in various levels of organisation. Insofar AI in particular is concerned, Benjamin argues that cases such as the AI-judged beauty content mentioned above or, to give another example, Twitter's cropping algorithm that seems to privilege lighter faces over darker ones, all account for the New Jim Code — that is, a very contemporary manifestation of a discriminatory architecture wherein racist assumptions are made intrinsic to the world's technological infrastructure. Benjamin's argument echoes concerns of Digital Colonialism and adds up to a body of work indicating that the technology upon which the contemporary world is run is alarmingly biased. To give a few more relevant examples, consider Noble (2018), discussing how search engines appear to systematically misrepresent black women with derogatory or overly sexualised images, or Daniels (2015), raising concerns of colour-blind racism throughout the way the internet manifests in the 21st century and arguing that racial inequality is a sad reality in the tech industry.

Under these premises, this project is part of a broader research endeavour taking a critical and de-colonialist stance on the current state of affairs in AI-driven image synthesis. It has surfaced critical reflection and prior experimentation with pre-configured AttnGAN models for image synthesis that would largely result in images that would still resemble western topography or architecture even when fed with descriptions of Eastern Mediterranean cultural affairs. In a fashion that is rather similar to the Feminist Data Set, this project revolves around the careful re-configuration of an existent AttnGAN model to arrive at a system that would genuinely account for non-western affairs.

3. Method

The employed pipeline builds upon the AttnGAN Fine-Grained T2I synthesis model introduced in Xu (2018). This is a multi-staged system, featuring a Text Encoder, an AttnGAN, a Deep Attentional Multimodal Similarity Model, and a CNN-driven image encoder. The Text Encoder is a bi-directional LSTM that maps the input to $R^{D \times T}$, D , being the dimension of the word vector T and the number of words. The last column \vec{e} is attached to the AttnGAN stage that, in this incarnation, features 3 hidden layers h_i that are sequentially inter-connected by means of an attentional model (a neural network):

$$F_i^{attn} : R^{D \times T} \times R^{\hat{D} \times N} \rightarrow R^{\hat{D} \times N}$$

$$e, h_{n-1} \mapsto h_n$$

Note that for h_0 , F^{ca} is used instead of, F^{attn} to convert \vec{e} into a conditional vector e that can be further propagated into the pipeline. Each column of h_i is a feature vector of a sub-region of the image so that F^{attn} returns a word-context vector to the next layer. Each h_i is also attached to an image generator (also a neural network); G_i these generators produce images of successively greater resolution, so that for each input two intermediate images are created alongside the resulting final one (generated by G_3). Fig. 1 illustrates the final and the intermediate images on a given input.

The final objective of the AttnGAN in Xu *et al* model is defined as

$$\mathcal{L} = \sum_0^i \mathcal{L}_{G_i} - \lambda \hat{\mathcal{L}}$$

where λ is a hyper-parameter to balance the two terms: the GAN loss as an approximation of conditional and unconditional distributions, and the word level fine-grained image-text matching loss $\hat{\mathcal{L}}$ that is computed as follows:


$$\hat{\mathcal{L}} = \mathcal{L}_1^w + \mathcal{L}_2^w + \mathcal{L}_1^s + \mathcal{L}_2^s$$

where \mathcal{L}_1^w and \mathcal{L}_2^w are the negative posterior log probability the images are matched with their corresponding text descriptions at word level and, symmetrically, *vice versa*; and where \mathcal{L}_1^s and \mathcal{L}_2^s the equivalent probability that images are matched with their descriptors at a sentence level and, symmetrically, *vice versa*. In other words, the model is designed to learn the attention model in a semi-supervised manner and with respect to the matching between images

and sentences (sequences of words), yet in a fine-grained fashion. Note that $\hat{\mathcal{L}}$ sees images as semantic vectors. It is, then, the Image Encoder's task to encode them appropriately. The latter is a CNN that builds upon Inception-v3 (Szegedy et al., 2016) to learn both global and local (concerning sub-regions of the original image) features.

The entire architecture is discussed in great detail in the original paper (Xu et al., 2018). Xu et al have trained the original model on the COCO and CUB datasets. The idea herein is to, instead, employ a DIY dataset pivoting on Cyprus and the broader region as seen by Baker. This is a decision of a rather conceptual character. Colonialism cannot be undone; so precisely because it is an integral part of the region's history, it has to be reflected on the training data. Data that should, nevertheless, still refer to this particular part of the world. Baker's book becomes, then, an invaluable resource precisely in that it does reflect a strictly colonialist view of Cyprus. Creating a custom dataset for AI training purposes is not trivial, however. Currently, the dataset is rather small in size, featuring hand-picked images from various sources that loosely concern Eastern-Mediterranean/Cypriot region that are labelled using excerpts from Baker's book. The dataset follows the structure of COCO in that the images are similarly complex and depicting a variety of objects and landscapes. In some detail, it comprises: (a) j-png images of landscapes, buildings, plants/trees, cityscapes, and similar; (b) a text description file for each image including five caption sentences; (c) a 'pickled' Python dictionary that maps keywords to numbers. In this way numeric lists can be generated that fully encode each image's textual descriptors.

Fig.1. First two rows illustrate the output of F^{lstm} and F^{attn} , respectively, indicating the five most attended words. The images on the last row are, from left to right, the 64×64 image generated by G_0 , the 128×128 generated by G_2 and the final 256×256 image output generated by G_3 . As is the case with the original AttnGAN pipeline, this system is also limited to 256×256 sized images, due to programming complexity (The text input is the same as the one in Fig. 5 below).



4. Results

Figures 3-7 illustrate some sample generated images (the final iteration) for a number of different text inputs passed through the pipeline. Figures 8-11 are given for reasons of comparison: these have been generated with the very same text inputs, yet this time employing the original (Xu et al., 2018) model without any modifications and trained on the COCO dataset. The difference between the two sets of images is striking, with the latter having nothing reminiscent of an Eastern Mediterranean region and been altogether different to the former. In order to better understand the results, the authors have conducted a formal, albeit short-scale, qualitative survey. A total of 27 people of various ages and socio-cultural backgrounds — as illustrated in figure 2 have responded to a set of questions about some generated images. The survey is structured in three parts: In parts (a) and (b) participants were asked to describe in their own words what they see in each image, their overall quality/aura and, for some images, the country or the geographic region they would associate them with. In part (c), they were given a text caption (the original input text used to generate the image) and they were asked to select the one, out of the three possible, image that best approximates/depicts it.

Fig.2. Survey demographics, note that participants were allowed to select more than one options insofar as their socio-cultural background is concerned.




Fig.3. Input text: *“The ruins of ancient cities offer no attraction to the traveller in this island, as nothing is to be seen upon the surface except disjointed stones and a few fallen columns of the commonest description”.*




Considering the image in figure 3 the vast majority of the respondents see some structure with ‘rocks’/‘stones’ (16 mentions) – possibly a ‘hill’, ‘mountain’ or a ‘cliff’ (12 mentions) in a ‘barren’/‘desert’ landscape (11 mentions); some specifically mention ‘archaeology’ (7 mentions). Then, 92.6% of participants did match this image correctly with the *“an archaeological site in Cyprus with disjointed stones”*. Responses about the overall aura/mood thereof vary; ‘sun’, ‘heat’, ‘age’, ‘calmness’, ‘haziness’, ‘surreal/strange’ qualities are mentioned, inter alia. To give a couple of example responses: *“peaceful atmosphere that can become intense or emotional”*; *“something surreal, something familiar”*. Most (at least 16) agreed that this image depicts an East-Mediterranean landscape, with 7 of them correctly identifying Cyprus.

There is a general consensus (18 mentions) that the image in figure 4 depicts ‘vegetation’, possibly with ‘yellow flowers’ (8 mentions); there are also some

references to the dryness of the landscape (6 mentions). Concerning the overall aura/mood here, responses oscillate between some positive/soothing atmosphere and a moody/uneasy feeling of discomfort. E.g.: *“dusk, like something is about to happen... but it doesn’t”*. Here, 74.1% of the participants correctly matched this image with *“trees crowded with oranges and lemons in a desolate land”*. At least 9 individuals see Eastern Mediterranean here, while another 5 see Northern Europe.

Fig. 4. Input text: *“Lemon and orange-trees of the largest size were crowded with fruit, and exhibited in the midst of a treeless and desolate country”*.



Several responses (8) also mention ‘vegetation’ on some ‘cliff’/ ‘mountain’/ ‘rocky terrain’ (6 mentions) in the case of the image in figure 5. The overall aura here is described as having an ‘abstract’/‘messy’/‘fractal’ quality (9 mentions) that is potentially ‘dangerous’/‘hellish’/‘intense’/‘violent’/‘psychedelic’/etc (16 mentions). Then, 59.3% of the participants correctly matched this image with *“Land covered with carob trees, whose dark green foliage gave a rich appearance to the shore”*, and most (12 individuals) correctly see an East Mediterranean landscape — with 6 of them pointing at Cyprus.

For the image in figure 6, the majority (7 mentions) of respondents sees ‘vegetation’ on a ‘dusty’/‘sandy’ landscape (9 mentions), possibly on a ‘mountain’ (4 mentions). Interestingly, many see water-related phenomena (7 mentions). There are also several responses suggesting a more abstract/fuzzy landscape, such as *“city on the clouds”*. This image also brought up unsettling feelings, with the vast majority of responses (14 mentions) concerning ‘mystery’, ‘gloom’, and *“something bad about to happen”*. Some participants described this image as

“Rough, dream-like, a real sense of distance”, or “unsettled, alien landscape”. Most (12 individuals) see Eastern Mediterranean here, while 6 of them correctly identify Cyprus. There are also some mentions to other Mediterranean regions (5 individuals) and U.S.A. (3 individuals). (This image was not featured in the third part of the questionnaire, so there are no data regarding how well it is matched to its generating text.)

Fig. 4. Input text: *“Land covered with carob trees, whose dark green foliage gave a rich appearance to the shore, broken by countless rocky bays and coves, filled with the waters of the Mediterranean. This was a lovely scene; I could not believe that I was in Cyprus—that whitey-brown-paper-coloured, desert, smitten, God-forsaken isle!”.*




Fig. 6. Input text: *“There are man stone columns lying useless among the heaps of ruins so common in Cyprus, that would form excellent rollers, but the idea of such an implement has never entered the Cypriote head.”*




Fig. 7. Input text: "Close to a cypress tree were rocks; steps led into a cave, where water dripped from the roof."



Fig. 8. Generated with Xu et al. (2018) model trained on COCO; input text same as in Fig. 3.



Fig. 9. Generated with Xu et al. (2018) model trained on COCO; input text is the same as in Fig. 4.



Fig. 10. Generated with Xu et al. (2018) model trained on COCO; input text is the same as in Fig. 5.



Fig. 11. Generated with Xu et al. (2018) model trained on COCO; input text is the same as in Fig. 6.



Fig. 12. Generated with Xu et al. (2018) model trained on COCO; input text is the same as in Fig. 7.



5. Discussion

Drawing upon the survey, it can be argued that to some significant extent the generated images (a) relate with the input text and sustain some traits found therein, and (b) somehow exhibit Cypriot or otherwise Eastern Mediterranean characteristics. Participants' responses vary significantly, of course, also depending on the image; notwithstanding, the above two conclusions are explicit in the results and safe to be considered largely granted. The situation is more convoluted when considering the overall aura/quality of the images. Still, there seem to be at least two distinct loci of convergence: (a) several responses suggest a rather uncanny, dream-like, and often unsettling/alien, quality insofar as all images are concerned — even if the exact wording is eclectic; (b) several other responses point at their abstractness/blurriness — this of no surprise of course since AI-generated imagery generally tends to be blurred. Note that these two characteristics do not contradict one another; imagery can be blurry/abstract and at the same time subjectively evocative/reminiscent of something. It is indeed often the case in the survey responses that the very same individual somehow refers to both qualities. To sum up, it can be said that the generated images appear somehow abstract and alien, yet sustaining qualities inherent to the Cypriot landscape and their generating text.

Can it, then, be argued that they also sustain traces of an explicitly biased view of Cyprus (as seen through Baker's gaze)? While there are some very strong indications pointing in that direction, this would be a largely unsupported claim to make based on the available data. The input text employed in all cases is taken from Baker's book, so that to find traces of the former in the resulting image is a very strong indication that Baker's gaze has left a prominent footprint therein. Even if the same text is used, the images in Fig. 8-12 are of a very different style so that, when trained with our database, the system does indeed acquire an overall different quality—one that could be thought of as reflecting a colonialist's gaze. A gaze, that nevertheless had a significant role in the island's modern history and in the formation of the national narrative. Yet, the possibility of dealing with a different phenomenon here, however feeble, cannot be excluded altogether—given, also, that there is not enough data concerning other non-western regions with which to compare the results. On the other hand, several of the participants' responses account for a 'weird', 'alien', and somehow exotic landscape in all the examples, which is exactly what Baker's view of Cyprus boils down into. Again, even if this is a strong indication pointing at this direction, the possibility that such qualities are merely congenital to the generating process cannot be proven false. So, while it cannot be argued

that Baker's colonialist look is specifically encoded in the pipeline, it can be supported that the latter does reflect the Cypriot geographical and socio-cultural reality to some significant extent. This may not sound as an achievement of some particular importance — in the end of the day it is not strange that the output of an AI architecture trained over some appropriate dataset is reminiscent of the former — to the authors' understanding, however, it is an important first step towards a conceptualisation process about post-colonial Cyprus and how it can claim presence in a digital reality that is being actively colonised at the moment.

From an ethical point of view, this project functions as a commentary on a much broader discourse: who creates vs who consumes the world's digital infrastructure? How can the non-western, second, or third world regions claim cultural presence in an ever-digitised world the governance of which is progressively trusted to intelligent technologies? What is the specific role of a small satellite and ex-colonised country such as Cyprus within such a context? These are very topical questions with deep ethical hues that concern all, but a few powerful, of the world's nations. Today, an ever-greater part of the world's population has online presence, while the AI-related technologies that are expected to shape the future of our societies are being developed and deployed by a very small group of powerful, and predominantly western, countries. When existent AI models already fail to account for anything that cannot be well appropriated with a Westerner's (white/male) perspective, and when AI research agenda is predominantly focused on performance and scalability (rather than ethics and equity), what is the future place of countries such as Cyprus in the world map of this emergent digital era?

To those claiming a multi-cultural digital era, the need to have off-centres somehow represented therein is a very pragmatic concern. Accordingly, this project is a concrete, albeit infinitesimal, step to speculate about how it may look like and to explore how it may be brought forth. What does it take to have traces of an individual (non-western) Cypriot cultural reality somehow still resonate through digital realms that have already been (digitally) colonised in many respects? And how would a non-colonised cyber-Cypriot culture look like? A lot of further research is necessary, of course, to even properly formulate questions of sort. Still, the authors aspire to contribute some valuable empirical data to a broader body of research that seeks to delineate the kind of tactics that are necessary for underrepresented territories/cultures to claim their genuine voice in a domain already ascribed with colonialist western ideals.

In tandem with the above, this endeavour ascribes to a broader critical and speculative design approach (Dunne 2008; 2013) where technology — AI programming in this particular case — is no longer thought of as the means to a solve some more or less well-defined problem. That is to say that whether the resulting imagery herein is indeed reminiscent of Cypriot landscapes/culture or not is not a question to be approached functionally and in a techno-solutionist fashion. The aim here is to rather arrive at a process and a series of artefacts that articulate some kind of critique and that echo a broader anti-colonialist ethos. In its essence, this is a political approach that is pursued in a hands-on research through (critical) design manner. The real methodology here lies in ‘hi-jacking’ and appropriating existent technology — one that generally reflects neo-colonialist digital ideals — to become a critical apparatus that reveals cracks in our understanding of everyday affairs. More importantly, to bring forth biases and inequalities, and, eventually, to speculate on alternative views of the world with more space for underrepresented and oppressed subjects.

6. Conclusion

It has been shown heretofore that it is indeed possible to appropriate an existent T2I GAN pipeline so as to arrive at results that properly reflect the local Cypriot reality — ones that, additionally, could be possibly shown to echo the region’s colonial past. Given that the employed model is almost identical to the one introduced in Xu et al. (2018) — a system originally designed to be an improvement over a series of other functional GANs — it is easy to see that in this particular case the possibility for not inherently western-biased imagery arises largely because a nonstandard DIY dataset is used. This is to have the focus shifted from how to improve performance, to how to derail it away to the implicit commands that are congenital to most readily available datasets. (An approach that is arguably similar in spirit to that of the “Feminist Data Set”.) Thus, the work hypothesis brought forth herein is rather simple: ‘hi-jacking’ some existent AI model, train it over some nonstandard dataset, and maybe employ it for purposes other than its original. In this fashion, they should turn into critical apparatuses that may potentially serve a broader (digital) de-colonialist agenda; both by means of expanding the polyphony of related research, and conceptually, bringing forth ethical and political questions.

Even if it is nontrivial to attach some model to a nonstandard, non-readily available, dataset, the approach looks and is straightforward. However, the importance of such an experiment lies not in the possibility of realistic imagery that meets some functional goal. But, rather, in that it concretely speculates on a

plausible cybernetic presence for the local Cypriot (post-colonial) society — thus, pragmatically suggesting a way out of the ongoing (digital) neo-colonialism. So to say, local affairs are empowered to generate upon their (AI-driven) cybernetic avatars and drawing upon local contexts, rather than upon what already offered by digital colonialism. In this vein, the herein described AttnGAN T2I pipeline aspires to become a de-colonialist critical apparatus — one that is not intended as the solution to some technological or design problem, but that rather raises ethical and political concerns.

Accordingly, the overall project is a very conscious attempt to critically engage with the possibility of alternative realities and research paradigms within AI-programming in general (with Cyprus as a point of reference). It is hoped that similar experiments therein, as well as in other off-centre regions, would eventually help produce a ‘toolkit’ of integrated tactics that could be applied to raise concerns of digital colonialism in AI research. This is a long-term affair, of course; nevertheless, this project is envisioned as contributing empirical data to this end: the output imagery and, most importantly, an account of a process that succeeds in giving some agency to local contexts in AI research.

7. Future Work

Future work primarily zooms in the production of a larger and much broader dataset that better reflects the local post-colonial reality. Manually doing so is labour-intensive, so that we also plan to trust algorithmic approaches, as well as ones pivoting on contributions by peers. Work could extend to include inherent biases within the software beyond the dataset. As a region of historical ethnic conflict, investigating possible biases within locally developed AI technologies and its impact on local communities is also considered for future work.

References

- Baker, Samuel White.**
1879. *Cyprus as I saw it in 1879*. London: Macmillan.
- Benjamin, Ruha.**
2019. Race after technology: Abolitionist tools for the new Jim Code. *Social Forces*, Volume 98, Issue 4, June 2020, Pages 1–3.
- Coleman, Daniel.**
2018. Digital colonialism: The 21st century scramble for Africa through the extraction and control of user data and the limitations of data protection laws. *Mich. J. Race & Law*. 24, 417.
- Daniels, Jessie.**
2015. “My brain database doesn’t see skin color” Color-blind racism in the technology industry and in theorizing the web. In: *American Behavioral Scientist* 59(11), 1377–1393.
- Dong, Yanlong, Ying Zhang, Lin Ma, Zhi Wang and Jiebo Luo.**
2020. *Unsupervised text-to-image synthesis*. Pattern Recognition p. 107573.
- Dunne, Anthony.**
2008. *Hertzian tales: Electronic products, aesthetic experience, and critical design*. Cambridge, MA: MIT press.
- Dunne, Anthony, and Fiona Raby.**
2013. *Speculative everything: design, fiction, and social dreaming*. Cambridge, MA: MIT press .
- Emir Ak, Kenan, Joo Hwee Lim, Jo Yew Tham and Ashraf Kassim.**
2019. Semantically consistent hierarchical text to fashion image synthesis with an enhanced-attentional generative adversarial network. In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*. pp. 0–0 .
- Goodfellow, Ian, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville and Yoshua Bengio.**
2014. Generative adversarial nets. In: *Advances in neural information processing systems*. pp. 2672–2680.
- Kwet, Michael.**
2019. Digital colonialism: US empire and the new imperialism in the global south. In: *Race & Class* 60(4), 3–26.
- Lin, Tsung-Yi, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár and C. Lawrence Zitnick.**
2014. Microsoft COCO: Common objects in context. In: *European conference on computer vision*. pp. 740–755. Springer.
- Nilsback, Maria-Elena and Andrew Zisserman.**
2008. Automated flower classification over a large number of classes. In: *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*. pp. 722–729. IEEE.
- Noble, Safiya Umoja.**
2018. *Algorithms of oppression: How search engines reinforce racism*. NYU Press.
- Pinto, Renata Avila.**
2019. Digital sovereignty or digital colonialism? In *International Journal on Human Rights*.
- Qiao, Tingting, Jing Zhang, Duanqing Xu and Dacheng Tao.**
2019. MirrorGAN: Learning text-to-image generation by redescription. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1505–1514.
- Reed, Scott, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele and Honglak Lee.**
2016. *Generative adversarial text to image synthesis*. arXiv preprint arXiv:1605.05396.
- Sinders, Caroline.**
2019. Making critical ethical software. In: *The Critical Makers Reader: (Un)learning Technology* p. 86. Amsterdam: Institute of Network Cultures.
- Szegedy, Christian, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens and Zbigniew Wojna.**
2016. Rethinking the inception architecture for computer vision. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2818–2826.
- Wah, C., Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie.**
2011. *The Caltech-UCSD birds-200-2011* dataset.
- Xu, Kelvin, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard Zemel and Yoshua Bengio.**
2015. Show Attend and Tell: *Neural image caption generation with visual attention*. arXiv Pre-Print 23.

Xu, Tao, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He.

2018. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1316–1324.

Young, Jason C.

2019. The new knowledge politics of digital colonialism. In: *Environment and Planning A: Economy and Space* 51(7), 1424–1441.

Zhang, Han, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang and Dimitris Metaxas.

2017. Stack-GAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 5907–5915.

Zhang, Han, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang and Dimitris Metaxas.

2018. Stack-GAN ++: Realistic image synthesis with stacked generative adversarial networks. In: *IEEE transactions on pattern analysis and machine intelligence* 41(8), 1947–1962.

Zhu, Minfeng, Pingbo Pan, Wei Chen and Yi Yang.

2019. Dm-GAN: Dynamic memory generative adversarial networks for text-to-image synthesis. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 5802–5810.